

**iVDGL/LHC Common Projects Facilities Meeting  
Brookhaven National Lab  
March 20-22, 2002**

**Meeting Summary**

The iVDGL Facilities and LHC Facilities Common Project Meeting was held at Brookhaven National Lab from March 20<sup>th</sup> to 22<sup>nd</sup>. This was the first face-to-face meeting of the iVDGL Facilities group, though the fourth meeting for the LHC Facilities Common Project Group.

The presentations from the meeting are available at  
<http://www.acf.bnl.gov/UserInfo/Events/iVDGL-Workshop.shtml>  
The following is a very short summary of the topics covered.

**March 20 Morning: Intro and Existing Facilities Overview**

Rich Baker gave a short introduction to and future directions of the committee.

All the participants were asked to give a short presentation about the computing facilities available at their sites. The committee has reasonable diversity of interests and philosophies, from very large, well staffed and managed computing clusters at the national labs to very small computing facilities located at Universities.

**March 20 Afternoon: Afternoon Disk Storage**

A presentation was made by Maurice Askinazi about the BNL disk storage system. The techniques for developing redundancy in data storage and data access were covered. Successes and troubles in recent BNL storage system acquisitions were also addressed.

Yujun Wu presented his work on disk performance evaluation. Yujun has compiled a suite of disk testing programs available from

<http://computing.fnal.gov/cms/disk>

Some disk performance numbers were presented, as was a first proposal for storage system validation.

**March 21 Morning: Linux Farms and Farm Management**

Tom Yanuklis presented the BNL Linux farm. Tom talked about the farm configuration, the split between reconstruction and analysis, and monitoring capabilities.

Mason Katz gave a presentation of the system configuration and monitoring capabilities of the Rocks cluster management software.

## **March 21 Afternoon: iVDGL**

The group had an afternoon brainstorming session about current activities and future directions. The first topic was a discussion about ways to present our computing centers to the grid not as facilities, but as services. Instead of letting high-level grid services access low level computing elements, like single computational nodes, descriptions of available computing resources will be published to the grid using monitoring and information tools. This would tend to put the interface to the global grid external to the inner workings of the local farm and give more control and flexibility to site administrators. The grid services can be broken into three areas: computation resources, storage resources, and network.

The second topic was a discussion of hardware testing in preparation for procurement. The group agreed to continue gathering experiment applications to use as benchmarks for evaluating systems. It was also agreed to use Yujun Wu's work with scripts for evaluating the performance of storage to establish common benchmarks for storage. The information and the performance numbers from both storage and CPU tests would be published to the iVDGL web pages. There was a discussion about what tests would be appropriate for local networking and it was decided, for the moment, not to conduct any.

## **March 22 Morning: Wrap up and Action Items**

### **Grid Enabled Facilities**

Over the next month: iVDGL Facilities members will give access to a limited number of computational nodes at their sites to designated contacts from other experiments. Site Contacts:

- UCSD Ian Fisk ([ifisk@ucsd.edu](mailto:ifisk@ucsd.edu))
- BNL Jason Smith ([smithj4@bnl.gov](mailto:smithj4@bnl.gov))
- Fermilab Hans Wenzel ([wenzel@fnal.gov](mailto:wenzel@fnal.gov))
- Florida Jorge Rodriguez ([Jorge@phys.ufl.edu](mailto:Jorge@phys.ufl.edu))
- Johns Hopkins Jan vandenBerg ([vincent@skysrv.pha.jhu.edu](mailto:vincent@skysrv.pha.jhu.edu))
- Caltech Suresh Singh ([suresh@cacr.caltech.edu](mailto:suresh@cacr.caltech.edu))

Over the next two months: we will attempt to establish software environments for the experiment code at "foreign" hardware facilities. Hopefully, non-privileged users can establish functional software environments. There is the possibility that some configuration will need to be performed by privileged users.

Site administrators will establish well-defined grid enabled batch queues for experiments using the globus gram. At the most basic, there will be an Atlas-queue, CMS-queue, LIGO-queue, and SDSS-queue, but it may be necessary to have additional queues for specific experiment tasks.

Dantong will design the initial schema for publishing site services using MDS.

Over the next six months: the experiments will attempt to configure grid enabled batch jobs to deploy the basic elements of the required software configurations from the batch job itself.

By SC2002 we expect to have demonstrated the use of grid enabled computing facilities outside the direct control of an experiment.

### **Hardware Procurement and Testing**

Until the Memorandums of Understanding are signed we do not expect any large-scale procurement of equipment using iVDGL funds.

An archived e-mail list called [facilities-dev@ivdgl.org](mailto:facilities-dev@ivdgl.org) will be created as a forum for hardware discussion.

In the next month, James Letts will put available experiment applications on the iVDGL web site. A non-expert user, Howard Brown, will test them. Sites with untested hardware configurations will test them and we will post results to the iVDGL web site. Yujun Wu will create a script, which uses his disk tools, so that basic disk performance quantities can be measured and recorded.

iVDGL Facilities will aim to schedule a face-to-face meeting this summer. We will investigate holding the meeting at one of the outreach sites; if there is one willing to host. Otherwise we will attempt to schedule the meeting at Fermilab.